Sparseness helps: Sparsity Augmented Collaborative Representation for Classification

Naveed Akhtar, Faisal Shafait, and Ajmal Mian

arXiv:1511.08956v1 [cs.CV] 29 Nov 2015

ABSTRACT

Many classification approaches first represent a test sample using the training samples of all the classes. This collaborative representation is then used to label the test sample. It was a common belief that sparseness of the representation is the key to success for this classification scheme. However, more recently, it has been claimed that it is the collaboration and not the sparseness that makes the scheme effective. This claim is attractive as it allows to relinquish the computationally expensive sparsity constraint over the representation. In this paper, we first extend the analysis supporting this claim and then show that sparseness explicitly contributes to improved classification, hence it should not be completely ignored for computational gains. Inspired by this result, we augment a dense collaborative representation with a sparse representation and propose an efficient classification method that capitalizes on the resulting representation. The augmented representation and the classification method work together meticulously to achieve higher accuracy and lower computational time compared to state-of-the-art collaborative representation based classification approaches. Experiments on benchmark face, object and action databases show the efficacy of our approach.

Index Terms—Multi-class classification, Sparse representation, Collaborative representation.

1 INTRODUCTION

SEVERAL recent approaches for multi-class classification (e.g. [4], [10], [28], [30], [31], [32]) exploit the representation of a test sample over a redundant basis, formed by the training samples (or their extracted features). This *collaborative representation* of the test sample, in which the training samples from different classes *collaborate* to approximate the test sample, is later used to decide its class label. Wright et al. [28] first demonstrated the impressive potential of this scheme for face recognition. Their approach additionally forces the representation to be sparse (i.e. it uses only a few vectors from the basis). Hence, it is called Sparse Representation based Classification (SRC).

- *N. Akhtar, F. Shafait and A. Mian are with the School of Computer Science and Software Engineering, The University of Western Australia, 35 Stirling Highway Crawley, 6009. WA.*

The success of SRC was followed up by its variants. For instance, Huang et al. [9] proposed a transformation-invariant SRC. Zhou et al. [33] combined Markov Random Fields with SRC for disguised faces. Similarly, Wagner et al. [26] enhanced SRC for the misalignment, pose and illumination invariant recognition. Effectiveness of these approaches also boosted significant research in dictionary learning [1] based multi-class classification [2], [10], [11], [19], [30]. Initially, the success of these approaches was attributed to the sparseness of the used representation. However, more recently, researchers have started questioning the role of sparsity in such approaches [20], [24], [31]. Among them, Zhang et al. [31] analyzed the working mechanism of SRC and claimed that it is the *collaboration and not the sparseness* of the representation that is the reason behind the effectiveness of SRC (and hence the related approaches). This result is rather widely acclaimed as it provides grounds to relinquish the computationally expensive sparsity constraint over the representation without sacrificing the classification accuracy.

In this paper, we first extend the analysis of Zhang et al. [31] and, in contrast to the original claim, we find that *sparseness of collaborative representation explicitly contributes to accurate classification*, hence it should not be completely ignored for computational gains. Motivated by this intuition, we propose a Sparsity Augmented Collaborative Representation based Classification scheme (SA-CRC) that uses both dense and sparse collaborative representations to decide the class label of a test sample. SA-CRC computes the dense representation using the regularized least squares method and greedily approximates the sparse representation using the Orthogonal Matching Pursuit (OMP) [25]. OMP's solution is used to augment the dense representation. Finally, the augmented representation is classified by capitalizing on its enriched discriminative properties. To that end, we propose an efficient classification method that avoids explicit computation of the reconstruction residuals for each class. We evaluate the proposed approach on two face databases [17], [8], one object category database [7] and a dataset for action recognition [21]. Extensive experiments show that our approach is not only more accurate than the state-of-the-art collaborative representation based classification approaches, its classification time is

also much lower than the approaches that ignore the sparsity altogether.

2 Problem Formulation

Let $\mathbf{\Phi} \in \mathbb{R}^{m \times N}$ denote the training data from C distinct classes, such that $\mathbf{\Phi} = [\mathbf{\Phi}_1, ..., \mathbf{\Phi}_i, ..., \mathbf{\Phi}_C]$. Each sub-matrix $\mathbf{\Phi}_i \in \mathbb{R}^{m \times n_i}$ pertains to a single class and $\sum_{i=1}^{C} n_i = N$. The columns of $\mathbf{\Phi}$ represent training samples, that are features extracted from images. Our goal is to develop an efficient multi-class classification scheme by collaboratively representing a test sample $\mathbf{y} \in \mathbb{R}^m$ over the training data[1]. A test sample is considered to be a feature vector that can be linearly approximated by the training samples. That is, $\mathbf{y} \approx \mathbf{\Phi}\boldsymbol{\alpha}$, where $\boldsymbol{\alpha} \in \mathbb{R}^N$ is the Collaborative Representation (CR) vector of the test sample. We allow $\mathbf{\Phi}$ to be a redundant set of basis vectors in $\mathbb{R}^m$. Furthermore, the subspaces spanned by the sub-matrices $\mathbf{\Phi}_{i \in \{1,...,C\}}$ are considered to be possibly overlapping, as this is often the case for the multi-class classification problems. Following the sparse representation literature [22], [27], we alternatively refer to $\mathbf{\Phi}$ as the *dictionary* and to its columns as the *dictionary atoms*. Furthermore, we generally refer to the *representation vector* (e.g. $\boldsymbol{\alpha}$) as *representation*, for brevity.

3 Related Work

Algorithm 1 CR-based Classification

Input: (a) Training data $\mathbf{\Phi}$, with samples normalized to have unit ℓ_2-norm. (b) Test sample $\mathbf{y}$. (c) Regularization parameter λ.

1: *Optimization:* Solve

$$\boldsymbol{\alpha} = \min_{\alpha} ||\mathbf{y} - \mathbf{\Phi}\boldsymbol{\alpha}||_2^2 + \lambda f(\boldsymbol{\alpha}), \tag{1}$$

where, $f(.)$ denotes a function and $||.||_p$ represents the ℓ_p-norm of a vector.

2: *Residual computation:* Compute class-specific reconstruction residuals $r_i(\mathbf{y}) = ||\mathbf{y} - \mathbf{\Phi}_i\acute{\boldsymbol{\alpha}}_i||_2, \ \forall i \in \{1, ..., C\}$, where $\acute{\boldsymbol{\alpha}}_i \in \mathbb{R}^{n_i}$ comprises the coefficients of $\boldsymbol{\alpha}$ corresponding to the i^{th} class.

3: *Labeling:* $\text{label}(\mathbf{y}) = \min_i\{r_i(\mathbf{y})\}$.

Output: $\text{label}(\mathbf{y})$.

Algorithm 1 presents the base-line scheme used by popular approaches (e.g. [28], [5], [31], [24], [11], [30]) that exploit collaborative representation in multi-class classification. The algorithm performs three key steps of (1) optimizing $\mathbf{y}$'s representation over a given dictionary, (2) computing class-specific reconstruction residuals $r_i(\mathbf{y})$, $\forall i \in \{1, ..., C\}$ and (3) labeling $\mathbf{y}$ using the computed residuals. In step (2), $\acute{\boldsymbol{\alpha}}_i \in \mathbb{R}^{n_i}$ comprises the coefficients of $\boldsymbol{\alpha}$ corresponding to the i^{th} class only. Hence, in step (3), $\mathbf{y}$ is assigned the label of the class that results in the smallest reconstruction residual. We can treat different approaches as special cases of the presented algorithm.

In SRC [28], $f(\boldsymbol{\alpha}) = ||\boldsymbol{\alpha}||_1$ in Eq. (1), which encourages the computed representation $\boldsymbol{\alpha}$ to be sparse. In Superposed-SRC (SSRC), Deng et al. [5] modified the residual computation step of SRC. For SSRC, $\mathbf{\Phi}$ consists of class centroids and sample-to-centroid differences. While computing the residuals, SSRC keeps the coefficients of $\boldsymbol{\alpha}$ corresponding to the sample-to-centroid differences fixed in each $\acute{\boldsymbol{\alpha}}_i$. The CR-based classifier proposed by Zhang et al. [31] uses $f(\boldsymbol{\alpha}) = ||\boldsymbol{\alpha}||_2$ and solves Eq. (1) using the Regularized Least Squares (RLS) method, hence denoted as CRC-RLS. Shi et al. [24] used $\lambda = 0$ in Eq. (1) and solved it as the standard least squares problem for face recognition. Chi and Porikli [3] used a linear combination of a CR-based classifier and a nearest subspace classifier [13] for improved classification performance.

Collaborative representation is also commonly used by discriminative dictionary learning techniques, e.g. [11], [30]. Although such approaches *learn* a dictionary instead of directly using the training data as $\mathbf{\Phi}$, explicit correspondence between the learned dictionary atoms and the class labels allows them to exploit the CR-based classification scheme. For instance, the Global Classifier (GC) used by Kong and Wang [11] is the same variant of Algorithm 1 that is used by SSRC [5]. The dictionary learned by the DL-COPAR algorithm [11] consists of COmmon atoms for all classes and PARticular atoms specific to each class. The particular atoms behave like class centroids whereas the common atoms act as centroid-to-sample differences in SSRC. Similarly, the GC used in the Fisher Discriminant Dictionary Learning (FDDL) [30] is a direct variant of CRC-RLS [31].

Another interesting direction of discriminative dictionary learning techniques, e.g. Label Consistent K-SVD (LC-KSVD) [10] and Discriminative K-SVD (D-KSVD) [32], is also related to CR-based classification. Such techniques learn collaborative dictionaries from the training data without enforcing strict correspondence between the class labels and the dictionary atoms. Due to the lack of such correspondence, the label of a test sample is chosen by maximizing a weighted sum of the coefficients of $\boldsymbol{\alpha}$, where the N-dimensional C weight-vectors are also learned during dictionary optimization. Among these weight-vectors, the i^{th} vector generally assigns large weights to the coefficients of $\boldsymbol{\alpha}$ corresponding to the dictionary atoms used commonly in representing the training data of the i^{th} class.

The above mentioned discriminative dictionary learning approaches classify a test sample using its representation over a collaborative set of features, learned directly from the training data. Therefore, they are considered to be instances of CR-based classification.

1. No explicit training of a machine learning algorithm is aimed, $\mathbf{\Phi}$ is conventionally referred as the *training* data [28], [31].

4 Collaboration and Sparsity

It is clear from Section 3 that many popular approaches directly exploit collaborative representation $\boldsymbol{\alpha}$ in classification. Whereas sparse representation based approaches (e.g. [28], [5]) associate the discriminative power of $\boldsymbol{\alpha}$ to its sparseness, there is an equal evidence in favor of discriminative abilities of dense representations [3], [24], [31]. In fact, it is also advocated that sparsity of the representation may not even be relevant to classification [20], [24], [31]. Zhang et al. [31] boosted the popularity of this notion by corroborating their claim with an analysis of the working mechanism of SRC. In Section 4.1, we closely follow this analysis to explain the role of collaboration in CR-based classification. We extend this analysis on the same lines of reasoning in Section 4.2 to show that collaboration alone is not sufficient for accurate classification. Section 4.3 discusses how sparseness additionally helps in this regard.

4.1 Why collaboration works?

We write the subspace spanned by the columns of $\boldsymbol{\Phi}$ as a set $\boldsymbol{\Psi}$. This subspace is geometrically illustrated as a plane in Fig. 1. Since a test sample $\mathbf{y}$ is approximated by the columns of $\boldsymbol{\Phi}$, we can write the approximation error as $\boldsymbol{\epsilon} = \mathbf{y} - \tilde{\mathbf{y}}$, where $\tilde{\mathbf{y}} = \boldsymbol{\Phi}\boldsymbol{\alpha} \subset \boldsymbol{\Psi}$ [2]. Let us represent the subspace spanned by the training data of the i^{th} class by a set $\boldsymbol{\Psi}_i$, where $\boldsymbol{\Psi}_i \subset \boldsymbol{\Psi}$. Without loss of generality, we can decompose $\tilde{\mathbf{y}}$ into two components, $\boldsymbol{\xi}_i$ and $\overline{\boldsymbol{\xi}_i}$ (illustrated in Fig. 1a) such that $\boldsymbol{\xi}_i \subset \boldsymbol{\Psi}_i$ and $\overline{\boldsymbol{\xi}_i} \subset \overline{\boldsymbol{\Psi}_i}$, where $\overline{\boldsymbol{\Psi}_i} = \bigcup_{j=1;j\neq i}^{C} \boldsymbol{\Psi}_j$. Similarly, the total approximation error $\boldsymbol{\epsilon}$ can itself be considered as a component of $\boldsymbol{\epsilon}_i$, where $||\boldsymbol{\epsilon}_i||_2$ represents the class-specific reconstruction residual $r_i(\mathbf{y})$, see step 2 of Algorithm 1.

To understand the working mechanism of CR-based classification, let $\mathbf{y}$ belong to the c^{th} class. In this case, $\tilde{\mathbf{y}} = \boldsymbol{\xi}_c + \overline{\boldsymbol{\xi}_c}$, i.e. $i = c$ in Fig. 1a. A CR-based classifier selects c as the label of $\mathbf{y}$ because $\boldsymbol{\epsilon}_i$ is expected to have the smallest length when $i = c$ [28], [31]. Zhang et al. [31] noted that this labeling criterion not only considers that the angle between $\tilde{\mathbf{y}}$ and $\boldsymbol{\xi}_c$ (i.e. β) is small, it also considers that the angle between $\boldsymbol{\xi}_c$ and $\overline{\boldsymbol{\xi}_c}$ (i.e. γ) is large. According to Zhang et al. [31], it is this *double-check with β and γ* (not the sparseness of the representation) that makes CR-based classification robust and effective. Therefore, they solved Eq. (1) using a computationally efficient regularized least squares method. The resulting dense collaborative representation was shown to be effective for face recognition, similar to sparse representation.

4.2 Why collaboration alone is not sufficient?

In the following text, we refer to a vector $\boldsymbol{\epsilon}_i$ as *class-specific error vector*. We present Lemma 4.1 regarding the underlying geometry of the class-specific error vectors involved in CR-based classification:

Lemma 4.1: For $i, j, k \in \{1, ..., C\}$, where $i \neq j \neq k$, the following holds: $\exists\ \boldsymbol{\epsilon}_i, \boldsymbol{\epsilon}_j$ such that $||\boldsymbol{\epsilon}_i||_2 = ||\boldsymbol{\epsilon}_j||_2$, while $\nexists\ \boldsymbol{\epsilon}_k$ such that $||\boldsymbol{\epsilon}_k||_2 < ||\boldsymbol{\epsilon}_i||_2$.

Proof: For our problem, the following holds under the law of sines, which can be verified from Fig. 1a:

$$\frac{||\tilde{\mathbf{y}}||_2}{\sin(\gamma)} = \frac{||\overline{\boldsymbol{\xi}_i}||_2}{\sin(\beta)}. \tag{2}$$

Also, $||\boldsymbol{\epsilon}_i||_2^2 = ||\boldsymbol{\epsilon}||_2^2 + ||\overline{\boldsymbol{\xi}_i}||_2^2$ because $\boldsymbol{\Psi} \perp \boldsymbol{\epsilon}$. From Eq. (2),

$$||\boldsymbol{\epsilon}_i||_2^2 = ||\boldsymbol{\epsilon}||_2^2 + \left(\frac{\sin(\beta)}{\sin(\gamma)}\right)^2 ||\tilde{\mathbf{y}}||_2^2. \tag{3}$$

Since $||\boldsymbol{\epsilon}||_2^2$ and $||\tilde{\mathbf{y}}||_2^2$ become constants once $\mathbf{y}$ is projected onto $\boldsymbol{\Psi}$, the condition that $\nexists\ \boldsymbol{\epsilon}_k$ s.t. $||\boldsymbol{\epsilon}_k||_2 < ||\boldsymbol{\epsilon}_i||_2$ holds when $\left(\sin(\beta)/\sin(\gamma)\right)^2$ is the minimum. However, for $\beta, \gamma \in [0, 2\pi]$ there is no unique minima for the given squared ratio. Hence, it is possible that $\exists\ \boldsymbol{\epsilon}_i, \boldsymbol{\epsilon}_j$ s.t. $||\boldsymbol{\epsilon}_i||_2 = ||\boldsymbol{\epsilon}_j||_2$, while $\nexists\ \boldsymbol{\epsilon}_k$ s.t. $||\boldsymbol{\epsilon}_k||_2 < ||\boldsymbol{\epsilon}_i||_2$.

Lemma 4.1, shows the possibility of existence of multiple class-specific error vectors with equal lengths when the length is minimized over the class labels. Figure 1b illustrates this possibility by drawing a circle of radius $||\overline{\boldsymbol{\xi}_i}||_2$ around point $\acute{o}$ on $\boldsymbol{\Psi}$. Any vector starting from a point on this circle (e.g. p, q) and ending at z will have the same length. For the labeling criterion of CR-based classification scheme, collaboration of the representation alone is not sufficient to indicate the best vector among these possible vectors. From Lemma 4.1, it is also evident that the double-check with β and γ mentioned by Zhang

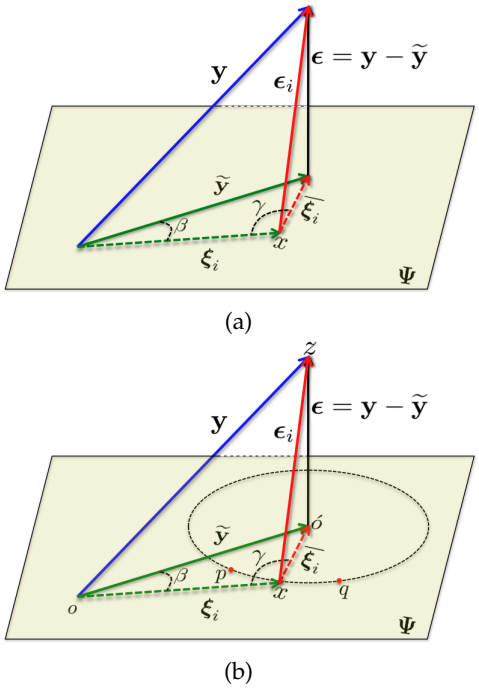

Fig. 1: Geometric illustration of the working mechanism of CR-based classification.

2. For $\boldsymbol{\Phi} \in \mathbb{R}^{m\times N}$, $\mathbf{y} \subset \boldsymbol{\Psi}$ when $N \to \infty$ and $\boldsymbol{\epsilon} \perp \tilde{\boldsymbol{\Psi}}$, where $\tilde{\boldsymbol{\Psi}} \subset \boldsymbol{\Psi}$. In that case, we are concerned with $\tilde{\boldsymbol{\Psi}}$ only, as $\mathbf{y}$ is considered to be approximated with a small error of bounded energy, i.e. $||\boldsymbol{\epsilon}||_2 \leqslant \varepsilon$. We exaggerate the error vector in figures for clarity.

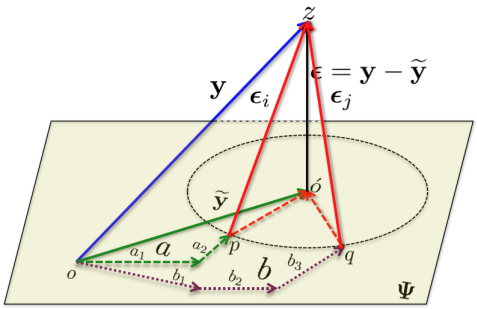

(a)

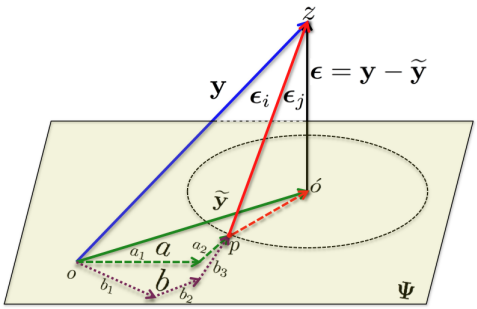

(b)

Fig. 2: Geometric illustration of the jointly exhaustive cases when $\exists \boldsymbol{\epsilon}_i, \boldsymbol{\epsilon}_j$ such that $||\boldsymbol{\epsilon}_i||_2 = ||\boldsymbol{\epsilon}_j||_2$: (a) $\boldsymbol{\xi}_i \neq \boldsymbol{\xi}_j$. (b) $\boldsymbol{\xi}_i = \boldsymbol{\xi}_j$. Here, $\boldsymbol{\xi}_i = \mathbf{a}$ and $\boldsymbol{\xi}_j = \mathbf{b}$ and the vectors are only displayed in terms of their components.

et al. [31] is essentially a *single-check* on the squared ratio of the sines of the angles. Thus, *CR-based classification without considering sparsity may not be as robust and effective as previously thought*.

4.3 How sparseness helps?

The above mentioned issue is inherent to CR-based classification scheme, with its roots in the redundancy in $\boldsymbol{\Phi}$. Simply computing a unique approximation of the representation, such as in CRC-RLS [31], does not resolve the issue because Lemma 4.1 still holds for the labeling step in Algorithm 1. To truly address the problem, a collaborative representation must be infused with additional information that finally results in using a suitable class-specific error vector in the labeling step. Sparsity constraint over the representation serves this purpose in CR-based classification.

To support our argument, in Fig. 2, we geometrically illustrate the two jointly exhaustive situations that can occur when two class-specific error vectors $\boldsymbol{\epsilon}_i$ and $\boldsymbol{\epsilon}_j$ have equal lengths, namely (a) $\boldsymbol{\xi}_i \neq \boldsymbol{\xi}_j$ and (b) $\boldsymbol{\xi}_i = \boldsymbol{\xi}_j$. In the figure, we denote $\boldsymbol{\xi}_i$ by $\mathbf{a}$ and $\boldsymbol{\xi}_j$ by $\mathbf{b}$ and show these vectors only by their components to avoid cluttering. In Fig. 2a, $\mathbf{a} \neq \mathbf{b}$ but $||\boldsymbol{\epsilon}_i||_2 = ||\boldsymbol{\epsilon}_j||_2$. In Fig. 2b, $\mathbf{a} = \mathbf{b} = \overrightarrow{op}$ and $||\boldsymbol{\epsilon}_i||_2 = ||\boldsymbol{\epsilon}_j||_2$. Although the class-specific residuals are equal in both cases, $\boldsymbol{\xi}_i$ and $\boldsymbol{\xi}_j$ can be distinguished based on their components. Intuitively, i (not j) represents the correct class of the test sample because $\boldsymbol{\xi}_i$ requires lesser number of components to produce the smallest class-specific residual. Fewer components of $\boldsymbol{\xi}_i$ implicates a sparser $\boldsymbol{\alpha}$. Hence, the sparsity constraint results in using a better class-specific error vector in the labeling step. Incidentally, the best performance of CR-based classification can be achieved by guaranteeing the representation to be the sparsest possible.

5 Proposed Approach

Computing the sparsest possible representation is generally NP-hard [18]. SRC [28] uses the ℓ_1-norm constraint to compute an approximate sparse representation, but the approach remains computationally expensive. On the other hand, computing a dense representation, such as in CRC-RLS [31], resolves the computational issues but it does not offer the advantages of sparsity. In the proposed classification scheme, we augment a dense representation with a greedily obtained approximate sparse representation. This augmentation enables accurate classification while keeping the approach computationally efficient.

Algorithm 2 presents the proposed scheme. In the first step, the algorithm optimizes two collaborative representations, i.e. $\breve{\boldsymbol{\alpha}}$ and $\hat{\boldsymbol{\alpha}}$. The dense representation $\breve{\boldsymbol{\alpha}}$ is computed using the regularized least squares method, whereas the sparse representation $\hat{\boldsymbol{\alpha}}$ is obtained by solving Eq. (4) using the Orthogonal Matching Pursuit (OMP) algorithm [25]. OMP iteratively selects k dictionary atoms to represent $\mathbf{y}$, hence, $\hat{\boldsymbol{\alpha}}$ has at most k non-zero coefficients, where k (the sparsity threshold) is determined by cross-validation. In each iteration, OMP chooses a dictionary atom by maximizing its correlation with an error vector. The error vector is computed as the difference between $\mathbf{y}$ and its orthogonal projection onto the subspace spanned by the already chosen atoms. For initialization, $\mathbf{y}$ itself is considered as the error.

As shown in step 2 of Algorithm 2, we add the sparse representation $\hat{\boldsymbol{\alpha}}$ to $\breve{\boldsymbol{\alpha}}$ and normalize the resulting vector to compute the augmented representation $\mathring{\boldsymbol{\alpha}}$. Despite being simple, this procedure greatly improves the discriminative abilities of the representation. We defer the discussion on the discriminative properties of $\mathring{\boldsymbol{\alpha}}$ to the upcoming paragraphs. These properties are exploited in step (3) of the algorithm to efficiently compute the label of the test sample $\mathbf{y}$. The labeling step uses a binary matrix $\mathbf{L} \in \mathbb{R}^{C \times N}$, that is provided as an input to the algorithm. For the i^{th} class, $\mathbf{L}$ contains n_i non-zero elements in its i^{th} row, at the indices corresponding to the columns of $\boldsymbol{\Phi}_i$. Thus, the i^{th} coefficient of $\mathbf{q} = \mathbf{L}\mathring{\boldsymbol{\alpha}}$ represents the sum of $\mathring{\boldsymbol{\alpha}}$'s coefficients corresponding to $\boldsymbol{\Phi}_i$. The label of the test sample is decided by maximizing the coefficients of $\mathbf{q}$. Empirical evidence for efficient and accurate classification using the proposed scheme is provided in Sections 6. Below, we analyze the reasons behind the improved performance of the approach.

For analysis, let us distribute the coefficient indices of a collaborative representation $\boldsymbol{\alpha}$ into two disjoint sets:

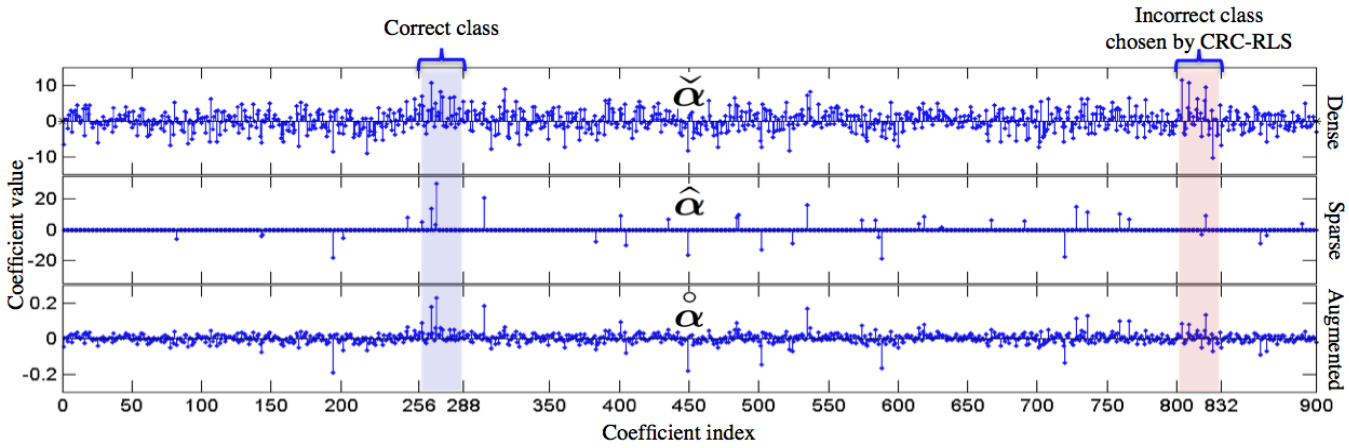

Fig. 3: Comparison of a test sample's representations for face recognition task using Extended YaleB database [8]: $\hat{\boldsymbol{\alpha}}$ consistently shows large positive values at the coefficients corresponding to the correct class. This culminates in correct classification by SA-CRC, whereas CRC-RLS [31] chooses incorrect class label for the test sample despite optimized parameters. Only first 900 coefficients are shown out of 1216.

Algorithm 2 Sparsity Augmented CR-based Classification

Input: (a) Training data $\boldsymbol{\Phi}$, with samples normalized in ℓ_2-norm. (b) Test sample $\mathbf{y}$. (c) Regularization parameter λ. (d) Sparsity threshold k. (e) Label matrix $\mathbf{L}$.

1: *Optimization:*
a) Compute $\check{\boldsymbol{\alpha}} = \mathbf{P}\mathbf{y}$, where, $\mathbf{P} = (\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi} + \lambda \mathbf{I}_N)^{-1}\boldsymbol{\Phi}^{\mathrm{T}}$.
b) Solve the following with greedy pursuit:

$$\hat{\boldsymbol{\alpha}} = \min_{\alpha} ||\mathbf{y} - \boldsymbol{\Phi}\boldsymbol{\alpha}||_2, \ s.t. \ ||\boldsymbol{\alpha}||_0 \leqslant k, \tag{4}$$

where, $||.||_0$ denotes the ℓ_0-pseudo norm.

2: *Augmentation:* Compute

$$\mathring{\boldsymbol{\alpha}} = \frac{\hat{\boldsymbol{\alpha}} + \check{\boldsymbol{\alpha}}}{||\hat{\boldsymbol{\alpha}} + \check{\boldsymbol{\alpha}}||_2} \tag{5}$$

3: *Labeling:* $\text{label}(\mathbf{y}) = \arg\max_i \{q_i\}$, where q_i denotes the i^{th} coefficient of $\mathbf{q} = \mathbf{L}\mathring{\boldsymbol{\alpha}}$.

Output: $\text{label}(\mathbf{y})$.

$\mathcal{A}_{\mathcal{H}} = \{i : \Xi_i > \delta\}$ and $\mathcal{A}_{\mathcal{L}} = \{j : \Xi_j \leqslant \delta\}$, where $\delta \in \mathbb{R}^+$ and $\Xi_n = \frac{\alpha_n^2}{||\boldsymbol{\alpha}||_2^2}$ with $\alpha_n \in \mathbb{R}$ denoting the n^{th} coefficient of $\boldsymbol{\alpha}$. The value Ξ_n represents the energy in the n^{th} coefficient, such that $\sum_{n=1}^{N} \Xi_n = 1$. If $\delta = 0$, $\mathcal{A}_{\mathcal{H}}$ contains the indices of non-zero coefficients of $\boldsymbol{\alpha}$, whereas $\mathcal{A}_{\mathcal{L}}$ comprises the indices of zero coefficients. Thus, the cardinality of the set $\mathcal{A}_{\mathcal{H}}$, i.e. $|\mathcal{A}_{\mathcal{H}}|$, defines the sparsity level of $\boldsymbol{\alpha}$. This remains true for $0 \leqslant \delta < \min_i \Xi_i$. Let $\boldsymbol{\alpha}^*$ denote the sparsest possible representation of $\mathbf{y}$ over $\boldsymbol{\Phi}$. We write the aforementioned sets for $\boldsymbol{\alpha}^*$ as $\mathcal{A}^*_{\mathcal{H}}$ and $\mathcal{A}^*_{\mathcal{L}}$. Furthermore, for any $\boldsymbol{\alpha}$, let us now fix $\delta = \left(\frac{\alpha^*_{min}}{||\boldsymbol{\alpha}^*||_2}\right)^2 - \varepsilon$, where α^*_{min} denotes the lowest energy coefficient of $\boldsymbol{\alpha}^*$. Hence, $|\mathcal{A}_{\mathcal{H}}|$ now counts the number of coefficients of $\boldsymbol{\alpha}$, each having at least the energy possessed by α^*_{min}. Therefore, henceforth, we refer to $|\mathcal{A}_{\mathcal{H}}|$ as the *effective sparsity* of the representation.

From Section 4, we know that $\boldsymbol{\alpha}^*$ is discriminative due to its sparsity. In practice, a representation $\mathring{\boldsymbol{\alpha}}$ is equally effective for classification if $|\mathcal{A}^{\circ}_{\mathcal{H}}| \approx |\mathcal{A}^*_{\mathcal{H}}|$ and the coefficients indexed in $\mathcal{A}^{\circ}_{\mathcal{H}}$ are discriminative[3]. For a dense representation $\check{\boldsymbol{\alpha}}$, $|\mathcal{A}^{\vee}_{\mathcal{H}}| \approx N \gg |\mathcal{A}^*_{\mathcal{H}}|$. Nevertheless, the representation is globally optimal. On the other hand, $|\mathcal{A}^{\wedge}_{\mathcal{H}}| \approx k \ll N$ for the sparse representation $\hat{\boldsymbol{\alpha}}$, but the representation is only locally optimal. However, $\hat{\boldsymbol{\alpha}}$ generally contains large positive coefficients at the indices corresponding to the correct class. For the other classes, most of the coefficients are either negative or have small positive values. This happens because OMP greedily assigns large values to the coefficients of $\hat{\boldsymbol{\alpha}}$ corresponding to the dictionary atoms that correlate more to $\mathbf{y}$, whereas $\mathbf{y}$ generally has a strong positive correlation with the samples of its own class. Thus, adding $\hat{\boldsymbol{\alpha}}$ to $\check{\boldsymbol{\alpha}}$ amplifies the coefficients of the correct class in the globally optimal solution. Figure 3 illustrate this phenomenon using an actual example of face recognition. In the figure, the coefficients of $\hat{\boldsymbol{\alpha}}$ are consistently positive and have relatively large values for the correct class. This finally results in dominant positive coefficients of $\mathring{\boldsymbol{\alpha}}$ for the correct class. For this example, CRC-RLS [31] is not able to identify the correct label of $\mathbf{y}$ despite optimized parameter settings, whereas the proposed approach classifies $\mathbf{y}$ correctly.

Notice that, the augmentation in Eq. 5 also results in $|\mathcal{A}^{\circ}_{\mathcal{H}}| \ll |\mathcal{A}^{\vee}_{\mathcal{H}}|$, because the procedure reduces the relative energy in the un-amplified coefficients of $\mathring{\boldsymbol{\alpha}}$. To illustrate the difference between the effective sparsity levels of the dense and the augmented representations, we plot the effective sparsity of the representations as a function of δ in Fig. 4. The plot is for actual face recognition task using Extended YaleB database [8]. The curve for the

3. We can safely ignore $\mathcal{A}^{\circ}_{\mathcal{L}}$ in this argument because the coefficients indexed in $\mathcal{A}^{\circ}_{\mathcal{L}}$ can be explicitly forced to zero, once $\mathcal{A}^{\circ}_{\mathcal{H}}$ is known.

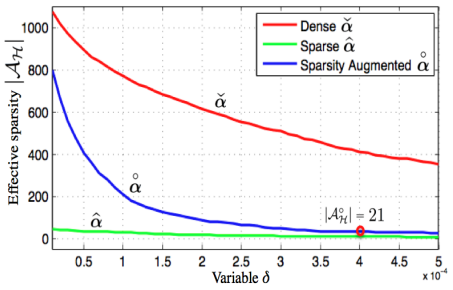

Fig. 4: Comparison of effective sparsity for face recognition, using Extended YaleB database [8].

augmented representation remains significantly lower than the curve for the dense representation. Moreover, for $\delta > 3 \times 10^{-4}$, $\mathring{\boldsymbol{\alpha}}$ is effectively almost as sparse as $\hat{\boldsymbol{\alpha}}$.

Considering the definition of effective sparsity, ideally, the coefficients of $\mathring{\boldsymbol{\alpha}}$ indexed in $\mathcal{A}^{\circ}_{\mathcal{L}}$ must be forced to zero before using the representation for classification. However, since δ is unknown, identifying the exact $\mathcal{A}^{\circ}_{\mathcal{L}}$ remains NP-hard. To resolve this issue, we design the labeling criterion that largely remains insensitive to the coefficients indexed in $\mathcal{A}^{\circ}_{\mathcal{L}}$. That is, instead of deciding the class label of a test sample based on the fidelity of its reconstruction, we directly integrate the coefficients of $\mathring{\boldsymbol{\alpha}}$ for each class separately. The largest integrated value indicates the correct class label. Due to the dominance of large values of the coefficients of the correct class in $\mathring{\boldsymbol{\alpha}}$, $\mathcal{A}^{\circ}_{\mathcal{L}}$ is not able to strongly influence the classification results. More precisely, our classification result remains as reliable as that obtained using an accurate representation with sparsity level $|\mathcal{A}^{\circ}_{\mathcal{H}}|$, under the mild worst-case condition $\sum_a - \sum_b > 2\sqrt{\delta}(n_b - n_a)$. Here, $\sum_a$ and $\sum_b$ denote the largest and the second largest integrated values of the coefficients, respectively, and n_b and n_a are the number of coefficients in $\mathring{\boldsymbol{\alpha}}$ contributing to $\sum_b$ and $\sum_a$ respectively, such that, each coefficient has energy less than δ. To exemplify, in Fig. 4, the classification results are as accurate as possible with sparsity level 21, unless $\sum_a - \sum_b \leqslant 0.04 \times (n_b - n_a)$. Typically, $\sum_a - \sum_b \in [0.1\ 0.3]$, whereas $n_b \approx n_a$. Since our labeling criterion does not need to compute reconstruction residuals for each class, we directly use the matrix $\mathbf{L}$ in step (3) of Algorithm 2. The matrix multiplication $\mathbf{L}\mathring{\boldsymbol{\alpha}}$ simultaneously integrates the coefficients for each class. Computationally, this makes our labeling step extremely effective.

6 Experiments

We evaluated the proposed approach using two face databases: AR database [17] and Extended YaleB [8], an object category database: Caltech-101 [7] and an action dataset: UCF sports actions [21]. These datasets are commonly used to benchmark the approaches that use collaborative representation for classification. We compare the performance of our approach to SRC [28], CRC-RLS [31], LC-KSVD [10], D-KSVD [32], FDDL [30] and DL-COPAR [11]. Unless mentioned otherwise, we performed our own experiments using the same training and testing partitions for all the approaches including the proposed approach. We carefully optimized the parameter values of the approaches using cross validation. For the existing techniques, these values are generally the same as those reported in the original works. However, for some cases, we used different values to favors these approaches. We explicitly mention these differences. For the dictionary learning approaches, the dictionaries are learned using the same training data that is directly used by SRC, CRC-RLS and the proposed approach.

We used the *author-provided codes* for CRC-RLS, LC-KSVD, FDDL and DL-COPAR. For SRC, we used the SPAMS toolbox [16] to solve the ℓ_1-norm minimization problem. For D-KSVD, we modified the public code of LC-KSVD [10]. In all the experiments, the proposed approach uses the implementation of OMP made public by Elad et al. [6]. The same implementation is used by LC-KSVD and D-KSVD. The proposed approach uses the sparsity threshold $k = 50$ for all the datasets. The regularization parameter λ is set to 0.003 for the face databases, 1.0 for the object database and 0.01 for the action database. Experiments are performed on an Intel Core i7-2600 CPU at 3.4 GHz with 8 GB RAM.

6.1 AR Database

The AR database [17] consists of over $4,000$ face images of 126 subjects. For each subject, 26 images are taken during two different sessions with large variations in terms of facial disguise, illumination and expressions. For our experiments, a 165×120 image was projected onto a 540-dimensional vector using a random projection matrix. Thus, the used samples are the Random-Face features [28]. We followed a common experimental protocol by selecting a subset of $2,600$ images of 50 male and 50 female subjects from the database. For each subject, 20 random images were chosen to create the training data and the remaining images were used for the test data.

In Table 1, we summarize the results on the AR database. The reported accuracies are the means (and standard deviations) of ten experiments. We also report the average time taken by each approach to classify a single test sample. For the parameter values of DL-COPAR, we followed the face recognition parameter settings in [11], which uses 15 atoms per class to represent class-specific data and 5 atoms to represent the commonalities. The Local Classifier [11] resulted in the best accuracy for DL-COPAR. For LC-KSVD [10] and D-KSVD [32] we set the sparsity threshold to 50 and the dictionary size to 1510 atoms for improved results. These values are different from the original works because these were found to give the best accuracies. For FDDL, we used the same parameter settings as [30] and the Global Classifier resulted in the best performance. For SRC [28], we set

TABLE 1: Recognition accuracies on the AR database [17] using Random-Face features. The reported average time (in milli-seconds) is for classifying a single test sample.

Method	Accuracy (%)	Time
DL-COPAR [11]	93.33 ± 1.69	40.01
LC-KSVD [10]	95.20 ± 1.22	1.56
D-KSVD [32]	95.41 ± 1.43	1.54
FDDL [30]	96.24 ± 1.01	51.23
SRC [28]	96.51 ± 1.36	69.91
CRC-RLS [31]	97.65 ± 0.67	4.46
SA-CRC (only RLS)	97.13 ± 0.74	0.07
SA-CRC (only OMP)	97.25 ± 0.43	2.00
SA-CRC (proposed)	$\mathbf{98.29 \pm 0.46}$	2.13

TABLE 2: Performance gain with SA-CRC in dictionary learning based multi-class classification. The average time (in milli-seconds) is for classifying a single test sample.

Method	Accuracy (%)	Time
K-SVD [1] + Lin. Classifier	94.06 ± 1.03	1.56
K-SVD $\mathbf{\Phi} \rightarrow$ SA-CRC	$\mathbf{95.65 \pm 0.66}$	1.61
ODL [15] + Lin. Classifier	94.60 ± 0.78	1.59
ODL $\mathbf{\Phi} \rightarrow$ SA-CRC	$\mathbf{95.33 \pm 0.68}$	1.62
LC-KSVD [10]	95.31 ± 1.06	1.55
LC-KSVD $\mathbf{\Phi}, \mathbf{L} \rightarrow$ SA-CRC	$\mathbf{96.44 \pm 0.99}$	1.61

the error tolerance $\varepsilon = 0.05$, as in the original work. For CRC-RLS [31], the regularization parameter λ is set to 0.003. This value is computed using the formula provided for λ for the face databases in [31]. Our cross-validation verified that this value results in the best performance of CRC-RLS.

Table 1 shows that the best results are achieved by the proposed approach, i.e. SA-CRC. We have also shown the results of our approach when we use only the Regularized Least Squares (RLS) or OMP in Algorithm 2. It is clear that using the augmented vector is better than using any of the two representation vectors alone. Notice that, due to the efficient classification criterion, our approach is much faster than CRC-RLS even when both OMP and RLS are used. The dictionaries used by LC-KSVD and D-KSVD are smaller in size as compared to the one used by SA-CRC, which results in slight computational advantage for these approaches. Nevertheless, accuracies of these approaches are much lower than SA-CRC.

In Table 2, we demonstrate the potential of SA-CRC for improving the performance of dictionary learning based multi-class classification approaches. The results are the mean values computed over ten experiments. We obtained the results in the first row as follows. First, K-SVD [1] is used to learn a dictionary containing 15 atoms per class. The sparse codes of the training data over the dictionary are used to compute a linear classifier, following [10]. A test sample is classified by first sparse coding it over the dictionary and then classifying its codes using the classifier. In the second row, we feed the same dictionary to SA-CRC as input $\mathbf{\Phi}$ instead (the test data remained the same). We also repeated the above procedure using the Online Dictionary Learning (ODL) approach [15] in place of K-SVD. The corresponding results are also reported. We can see a clear gain in the classification accuracies using SA-CRC in both cases. The table also compares the classification performance of LC-KSVD [10] with its enhancement using SA-CRC. For the enhancement, we replaced the classification stage of LC-KSVD with SA-CRC. That is, the dictionary and the weight matrix learned by LC-KSVD are directly used as SA-CRC's inputs $\mathbf{\Phi}$ and $\mathbf{L}$, respectively. There is a clear improvement in the classification performance of LC-KSVD after this modification. The performance of other discriminative dictionary learning approaches can also be improved using SA-CRC. Results in Table 1 and 2 demonstrate the potential of sparsity augmented collaborative representation for improved CR-based classification across the board.

TABLE 3: Recognition accuracies with Random-Face features on the Extended YaleB database [8]. The average time (in milli-seconds) is for classifying a single sample.

Method	Accuracy (%)	Time
D-KSVD [32]	94.71 ± 0.45	0.41
DL-COPAR [11]	94.87 ± 0.55	31.75
LC-KSVD [10]	95.38 ± 0.64	0.42
FDDL [30]	96.19 ± 0.71	58.19
SRC [28]	97.06 ± 0.41	68.12
CRC-RLS [31]	97.81 ± 0.44	2.41
SA-CRC	$\mathbf{98.32 \pm 0.43}$	1.23

6.2 Extended YaleB

The Extended YaleB face database [8] comprises $2,414$ images of 38 subjects. Each subject has about 64 samples acquired under varying illumination conditions with different expressions. For this database, 192×168 cropped images were projected onto a 504-dimensional vector to obtain the Random-Face features. For evaluation, we used a common experimental setting, where half of the available features of each subject were used for the training data and the remaining half were used in testing.

In Table 3, we show the results on Extended YaleB. For D-KSVD and LC-KSVD we used 600 dictionary atoms as they gave the best accuracies. The remaining parameters of these algorithms were set to the original values reported in [10]. We set the regularization parameter of CRC-RLS to 0.002 for this database, as guided by [31] and dictated by cross-validation. For the remaining approaches, the parameter values reported in Section 6.1 also resulted in their best performances for this database, hence they were kept the same. Again, SA-CRC is able to outperform the existing techniques. Although SA-CRC attains only a slight advantage over CRC-RLS in terms of accuracy for this dataset, it is able to classify a test sample *almost twice* as fast as CRC-RLS.

6.3 Caltech-101

Caltech-101 database [7] contains $9,144$ images from 101 object classes and one class of background images. The

TABLE 4: Classification accuracies (%) on the Caltech-101 dataset [7] using the spatial pyramid features.

Training samples	10	15	20	25	30
SRC [28]	57.8	63.3	67.2	69.2	71.8
CRC-RLS [31]	59.4	64.8	68.0	69.3	71.8
DL-COPAR [11]	58.4	65.1	69.3	71.1	72.5
FDDL [30]	59.7	66.6	69.1	71.3	72.9
D-KSVD [32]	60.7	66.3	69.6	71.0	73.1
LC-KSVD [10]	62.9	67.3	70.3	72.6	73.4
SA-CRC	**63.2**	**68.2**	**71.9**	**73.6**	**76.1**

TABLE 5: Computation time (in seconds) for classification on Caltech-101 dataset [7].

Method	Time	Method	Time
D-KSVD [32]	19.80	SA-CRC	21.43
LC-KSVD [10]	19.91	CRC-RLS [31]	130.41

classes include diverse categories of object (e.g. trees, minarets, signs) with significant shape variation within a category. For each class, the number of available images vary between 31 and 800. In our experiments, the used image feature descriptors were obtained by the following procedure. First, the SIFT descriptors [14] were extracted from 16×16 patches. Based on these descriptors, spatial pyramid features [12] were extracted with 1×1, 2×2 and 4×4 grids. For extracting these features, the codebook was trained using k-means, where $k = 1024$. Finally, the dimension of a feature was reduced to $3,000$ using PCA. Following a common experimental setting, we created 5 sets of train and test data with the extracted features. These sets consisted of $10, 15, 20, 25$ and 30 training samples per class, whereas the remaining samples were used as the test data in each case. We repeated our experiments ten times, every time selecting the training and testing data randomly.

Table 4 shows the mean classification accuracies for our experiments. We used the error tolerance of 10^{-6} for SRC, which gave the best results. The regularization parameter $\lambda = 1.0$ for CRC-RLS. The same value of the regularization parameter is used by SA-CRC to solve the RLS problem. For FDDL, we used the parameter settings of the object categorization experiments in [30]. DL-COPAR and LC-KSVD use the same settings as in the original works for the same database. These settings also resulted in their best performance for our data. For D-KSVD, the settings used by [10] showed the best results.

It is clear from Table 4 that SA-CRC consistently outperforms the existing approaches. In Table 5, we also report the classification time (for the complete test data) of the four most efficient approaches. The time is computed when 30 samples were used for training and the rest were used for testing. We can see that SA-CRC is more than six times faster than CRC-RLS and its timings are comparable to those of the efficient discriminative dictionary learning approaches. Note that, D-KSVD and LC-KSVD also required around 90 minutes of training.

TABLE 6: Classification accuracies (%) on UCF Sports Action dataset [21] using the action bank features.

Method	Acc.	Method	Acc.
Sadanand [23]	90.7	FDDL [30]	93.6
DL-COPAR [11]	90.7	LC-KSVD [10]	94.2
D-KSVD [32]	93.4	CRC-RLS [31]	94.4
SRC [28]	93.5	SA-CRC	**95.7**

6.4 UCF Sports Actions

The UCF Sports Action dataset [21] contains video sequences collected from different broadcast sports channels. The videos are from 10 categories of sports actions (e.g. diving, lifting, running). For this dataset, we directly used the action bank features made public by Sadanand and Corso [23]. A common evaluation protocol was followed in our experiments, where a fivefold cross validation was performed using four folds for training and the remaining one for testing. The results in Table 6 are the average accuracies of the five experiments. The reported accuracies of Sadanand [23], DL-COPAR and FDDL are taken directly from [29], where the same experimental protocol has been followed. Our parameter optimization for FDDL and DL-COPAR could not achieve these accuracies. For the remaining approaches the results are reported on the same folds using the optimized parameter values. For SRC, the error tolerance was set to 10^{-6} and 50 dictionary atoms were used for LC-KSVD and D-KSVD. The same number of atoms were used by Jiang et al. [10]. We used $\lambda = 0.01$ for both CRC-RLS and SA-CRC, which resulted in their best performance. For all the five experiments combined, the classification time was 0.04 seconds for SA-CRC and 0.31 seconds for CRC-RLS.

7 Discussion

Our approach requires a regularization parameter λ and sparsity threshold k as the input parameters for a given pair of $\mathbf{\Phi}$ and its label matrix $\mathbf{L}$. In our experiments, we optimized the values of these parameters by cross-vlaidation using the following systematic procedure. First, λ was optimized by executing Algorithm 2 without step 1(b) and considering $\hat{\boldsymbol{\alpha}}$ to be a zero vector in Eq. 5. Then k was optimized by fixing λ to the optimized value and executing the complete Algorithm. The parameters were further fine-tuned to nearby values when doing so yielded better performance.

To show the behavior of SA-CRC for different parameter values, in Fig. 5, we plot the classification accuracy of SA-CRC as a function of λ and k by fixing one parameter and varying the other. We also include results of CRC-RLS [31] for comparison. Plots in Fig. 5a, are for AR database [17] where we followed the experimental protocol of [31]. In the first plot (from left), we fixed k to 50 and varied λ. Clearly, SA-CRC consistently outperforms CRC-RLS and the results are less sensitive to the values of λ once k is fixed to an optimized value. In the second

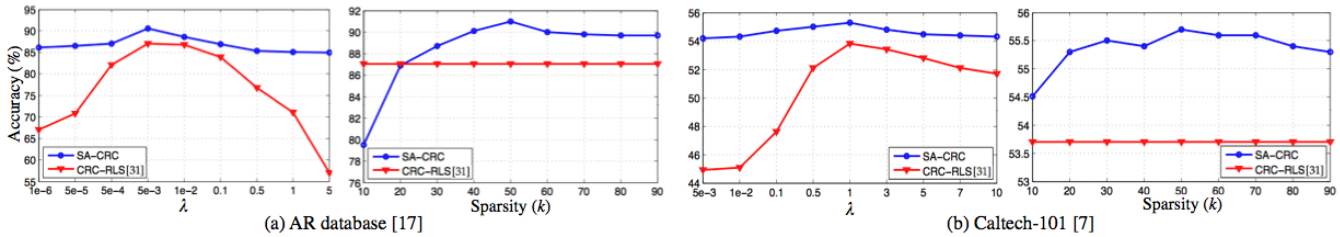

Fig. 5: Accuracy as a function of parameters: One parameter is fixed for SA-CRC while the other is varied. Both CRC-RLS [31] and SA-CRC always use the same values of λ. (a) k is fixed to 50 in the first plot (from left) and λ is fixed to 0.003 in the second. (b) k is fixed to 50 and λ to 1, respectively.

plot, we used $\lambda = 0.003$ for both SA-CRC and CRC-RLS and varied k for SA-CRC. Again, for $k > 20$, SA-CRC consistently outperforms CRC-RLS. Qualitatively speaking, Fig. 5a shows a typical relationship between the performance of CRC-RLS and SA-CRC that was observed in our experiments on face databases.

In Fig. 5b, we repeated the same experiment for the object dataset, Caltech-101 [7]. To fix the parameter values, we used $k = 50$ and $\lambda = 1$. For this experiment, we used five samples per class for training and the rest for testing. Again, the obtained results consistently favor SA-CRC. Qualitatively similar behavior was observed for all the train/test partitions used in Table 4.

8 Conclusion

In contrast to a popular existing notion, we showed that sparsity of a Collaborative Representation (CR) plays an explicit role in accurate CR-based classification, hence it should not be completely ignored for computational gains. Inspired by this result, we propose a Sparsity Augmented Collaborative Representation based Classification scheme (SA-CRC) that augments a dense collaborative representation with an efficiently computed sparse representation. The resulting representation is classified using a efficient method. Extensive experiments for face, action and object classification establish the effectiveness of SA-CRC in terms of accuracy as well as computational efficiency.

Acknowledgments

This work was supported by ARC Grants DP110102399 and DP1096801.

References

[1] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE TSP*, 54(11):4311–4322, Nov. 2006.

[2] N. Akhtar, F. Shafait, and A. Mian. Discriminative Bayesian Dictionary Learning for Classification. *arXiv preprint arXiv:1503.07989*, 2015.

[3] Y. Chi and F. Porikli. Classification and boosting with multiple collaborative representations. *IEEE TPAMI*, 36(8):1519 – 1531, August 2014.

[4] W. Deng, J. Hu, and J. Guo. Extended SRC: Undersampled face recognition via intraclass variant dictionary. *IEEE TPAMI*, 34(9):1864–1870, Sept 2012.

[5] W. Deng, J. Hu, and J. Guo. In defense of sparsity based face recognition. In *CVPR*, 2013.

[6] M. Elad, R. Rubinstein, and M. Zibulevsky. Efficient implementation of the K-SVD algorithm using Batch Orthogonal Matching Pursuit. Technical report, CS, Technion, 2008.

[7] L. Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training examples: an incremental Bayesian approach tested on 101 object categories. In *IEEE CVPR 2004, Workshop on Generative-Model Based Vision*, 2004.

[8] A. Georghiades, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE TPAMI*, 23(6):643–660, 2001.

[9] J. Huang, X. Huang, and D. Metaxas. Simultaneous image transformation and sparse representation recovery. In *CVPR*, 2008.

[10] Z. Jiang, Z. Lin, and L. Davis. Label consistent K-SVD: Learning a discriminative dictionary for recognition. *IEEE TPAMI*, 35(11):2651–2664, Nov 2013.

[11] S. Kong and D. Wang. A dictionary learning approach for classification: Separating the particularity and the commonality. In *ECCV*, 2012.

[12] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, 2006.

[13] K.-C. Lee, J. Ho, and D. Kriegman. Acquiring linear subspaces for face recognition under variable lighting. *IEEE TPAMI*, 27(5):684–698, May 2005.

[14] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, Nov. 2004.

[15] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary larning for sparse coding. In *ICML*, 2009.

[16] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, Mar. 2010.

[17] A. M. Martinez and R. Benavente. The AR Face Database. Technical report, CVC, June 1998.

[18] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM J. Comput.*, 24(2):227–234, Apr. 1995.

[19] I. Ramirez, P. Sprechmann, and G. Sapiro. Classification and clustering via dictionary learning with structured incoherence and shared features. In *CVPR*, 2010.

[20] R. Rigamonti, M. Brown, and V. Lepetit. Are sparse representations really relevant for image classification? In *CVPR*, 2011.

[21] M. Rodriguez, J. Ahmed, and M. Shah. Action mach a spatio-temporal maximum average correlation height filter for action recognition. In *CVPR*, 2008.

[22] R. Rubinstein, A. Bruckstein, and M. Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045–1057, June 2010.

[23] S. Sadanand and J. J. Corso. Action bank: A high-level representation of activity in video. In *CVPR*, 2012.

[24] Q. Shi, A. Eriksson, A. van den Hengel, and C. Shen. Is face recognition really a compressive sensing problem? In *CVPR*, 2011.

[25] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE TIT*, 53(12):4655–4666, 2007.

[26] A. Wagner, J. Wright, A. Ganesh, Z. Zhou, and Y. Ma. Towards a practical face recognition system: Robust registration and illumination by sparse representation. In *CVPR*, 2009.

[27] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. S. Huang, and S. Yan. Sparse Representation for Computer Vision and Pattern Recognition. *Proceedings of the IEEE*, 98(6):1031–1044, June 2010.

[28] J. Wright, A. Y. Yang, A. Ganesh, S. S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE TPAMI*, 31(2):210–227, Feb. 2009.

[29] M. Yang, D. Dai, L. Shen, and L. Van Gool. Latent dictionary learning for sparse representation based classification. In *CVPR*, 2014.
[30] M. Yang, L. Zhang, X. Feng, and D. Zhang. Sparse representation based fisher discrimination dictionary learning for image classification. *IJCV*, 109(3):209–232, 2014.
[31] L. Zhang, M. Yang, and X. Feng. Sparse representation or collaborative representation: Which helps face recognition? In *ICCV*, 2011.
[32] Q. Zhang and B. Li. Discriminative k-svd for dictionary learning in face recognition. In *CVPR*, 2010.
[33] Z. Zhou, A. Wagner, H. Mobahi, J. Wright, and Y. Ma. Face recognition with contiguous occlusion using markov random fields. In *ICCV*, 2009.